Submitted to
manuscript 2015

Authors are encouraged to submit new papers to INFORMS journals by means of a style file template, which includes the journal title. However, use of a template does not certify that the paper has been accepted for publication in the named journal. INFORMS journal templates are for the exclusive purpose of submitting to an INFORMS journal and should not be used to distribute the papers in print or online or to submit the papers to another publication.

arXiv:1605.03205v1 [cs.SI] 10 May 2016

Profit-aware Team Grouping in Social Networks: A Generalized Cover Decomposition Approach

Shaojie Tang
Naveen Jindal School of Management, The University of Texas at Dallas

In this paper, we investigate the profit-aware team grouping problem in social networks. We consider a setting in which people possess different skills and compatibility among these individuals is captured by a social network. Here, we assume a collection of tasks, where each task requires a specific set of skills, and yields a different profit upon completion. Active and qualified individuals may collaborate with each other in the form of *teams* to accomplish a set of tasks. Our goal is to find a grouping method that maximizes the total profit of the tasks that these teams can complete. Any feasible grouping must satisfy the following three conditions: (i) each team possesses all skills required by the task, (ii) individuals within the same team are social compatible, and (iii) each individual is not overloaded. We refer to this as the TeamGrouping problem. Our work presents a detailed analysis of the computational complexity of the problem, and propose a LP-based approximation algorithm to tackle it and its variants. Although we focus on team grouping in this paper, our results apply to a broad range of optimization problems that can be formulated as a cover decomposition problem.

Key words: approximation algorithm; team formation; cover decomposition

1. Introduction

In this paper, we consider the problem of grouping teams on a networked community of people with diverse skill sets. We consider a setting in which people possess different skills and compatibility among these individuals is captured by a social network. Here, we assume a collection of tasks, where each task requires a specific set of skills, and yields a different profit upon completion. Active and qualified individuals may collaborate with each other in the form of *teams* to accomplish a set of tasks. Our goal is to find a grouping method that maximizes the total profit of the tasks that these teams can complete. One relevant example is from the domain of online labor markets, such as Freelancer

(`www.freelancer.com`), Upwork (`www.upwork.com`), and Guru (`www.guru.com`). In these online platforms, freelancers with various skills can be hired to work on different types of projects. Instead of just working independently, more and more freelancers are realizing that it is more beneficial to work as a team, together with other solo freelancers who have complementary skills (Golshan et al. 2014). This allows them to expand their talent pool and achieve better load balance. Nowadays many major platforms in this area such as Upwork has provided team-hiring services to their enterprise customers.

In this paper, we formalize the profit-aware team grouping problem as follows: we assume a set of m individuals $\mathcal{V}$ and a set of n skills $\mathcal{S}$. Each individual $u \in \mathcal{V}$, is represented by a subset of skills, i.e., $u \subseteq S$; these are the skills that the individual possesses. We also assume a set of tasks $\mathcal{T}$, every task $t \in \mathcal{T}$ can also be represented by the set of skills that are required in order for the task to be completed (i.e., $t \subseteq \mathcal{S}$). Finally, every task t is associated with a profit λ_t, this could be the benefit that the completion of a task will yield for the platform. The team grouping problem is to group individuals to different teams and assign them to different tasks satisfying the following three conditions: (i) each team possesses all skills required by the task, (ii) individuals within the same team have high social compatibility, and (iii) each individual is not overloaded. Our goal is to maximize the sum of profits from all tasks that can be performed. We refer to this as the TeamGrouping problem. It was worth noting that the social compatibility among individuals can be interpreted in many ways. In this work, we model the *social compatibility* by means of a social network. One natural and popular option with respect to capturing the underlying social compatibility of a team is *connectivity*. This follows the approach of (Lappas et al. 2009) and requires that each team forms a connected graph. Other options to measure social compatibility include the diameter of a team (Anagnostopoulos et al. 2012), i.e., the induced graph of any team in G must have small diameter. Fortunately, our results are not restricted to any specific measures of social compatibility, instead, we propose a general framework that works for any reasonable measure.

For example, assume there are three IT projects requiring different skills: the first task has profit $\lambda_1 = \$50$ requiring skills $t_1 = \{\text{HTML, MySQL, JavaScript, PHP}\}$, the second task has profit $\lambda_2 = \$10$ requiring skills $t_2 = \{\text{JavaScript, HTML}\}$, and the last task has profit $\lambda_3 = \$5$ requiring skills $t_3 = \{\text{PHP}\}$. Also assume there are three individuals, $\{a, b, c\}$ with the following backgrounds: $a = \{\text{HTML, MySQL}\}$, $b = \{\text{JavaScript}\}$, $c =$

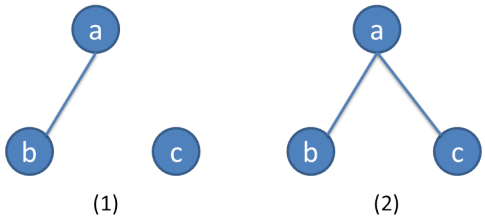

Figure 1 Two social networks.

{HTML, PHP}. Our basic formulation requires that each individual can only participate in one team, and all team members must be connected. We consider two different social networks as represented in Fig. 1. The most profitable grouping in Fig. 1 (1) is to assign team $\{a, b\}$ to t_2, and team $\{c\}$ to t_3, which yields \$15 in profit. This is because a and b are connected, but c is isolated. Consider another social network in Fig. 1 (2) by contrast, since the induced graph of all three individuals is connected, the most profitable grouping is to assign team $\{a, b, c\}$ to t_1, which yields \$50 in profit.

Contributions: To the best of our knowledge we are the first to define and study the TeamGrouping problem and its variants. We summarize our contributions as follows:

- We show that this problem is $1/\ln m$ hard to approximate, i.e., it is NP-Hard to find a solution with approximation ratio larger than $1/\ln m$.
- We propose a LP-based algorithm with approximation ratio $\max\{\mu/\Delta, \mu/2\sqrt{m}\}$ where Δ denotes the size of the largest minimal team and μ is the approximation ratio of MincostTeamSelection problem. If there is no constraint on social compatibility, this ratio is equivalent to $\max\{\ln n/n, \ln n/2\sqrt{m}\}$.
- We also consider two extensions of the basic model. In the first extension, we relax the assumption that each individual can only participate in one task by allowing individuals to have different load limits. In the second part, we consider the scenario when each task can only be performed by a fixed number of teams or times. We develop effective approximation algorithms to tackle both extensions.
- Although we restrict our attention to the profit-aware team grouping problem in this paper, our results apply to other applications such as lifetime maximization problem in wireless networks (Bagaria et al. 2013), resource allocation and scheduling problems (Pananjady et al. 2014), and supply chain management problems (Lu 2011). In this sense, this research contributes fundamentally to the development of approximate solutions for any problems that fall into the family of generalized cover decomposition problem.

The remaining of this paper is organized as follows. In Section 2, we review the literature on team formation and disjoint set cover. We introduce the formulation of our problem in Section 3. In Section 4, we present our LP-based approximation algorithms. We also study two extensions of the basic model in Section 5. We summarize this paper in Section 6. Most notations used in this paper are summarized in Table 1.

Table 1 Symbol table.

Notation	Meaning
n, m, k	number of skills, individuals, tasks
Δ	the size of the largest minimal team
C_{ti}	the i-th team that covers task t
μ	approximation ratio of MincostTeamSelection problem
x^*_{ti}	(approximate) solution of primal LP
$\mathcal{N}(C_{ti})$	C_{ti}'s adjacent teams
$\mathcal{C}_t$	set of teams that cover task t
$\mathbf{C}$	$\mathbf{C} = \{\mathcal{C}_1, \cdots, \mathcal{C}_k\}$
$\mathcal{C}^H_t$	set of teams on task t corresponding to the separating hyper-planes
$\mathbf{C}^H$	$\mathbf{C}^H = \{\mathcal{C}^H_1, \cdots, \mathcal{C}^H_k\}$
$\mathbf{C}^I$	set of input teams of LP rounding

2. Related Work

To the best of our knowledge we are the first to formulate and study the team grouping problem and its variants. However, our work is closely related to other team formation and cluster hiring problems. Lappas et al. (2009) introduce the minimum cost team formation problem. Given a set of skills that need to be covered and social network, their objective is to select a team of experts that can cover all required skills, while ensuring efficient communication between the team's members. There is a considerable amount of literature on this topic and its variants (Kargar et al. 2013, Dorn and Dustdar 2010, Gajewar and Sarma 2012, Kargar and An 2011, Li and Shan 2010, Sozio and Gionis 2010). In (Golshan et al. 2014), they study cluster hire problem, where the objective is to hire a profit-maximizing team of experts with the ability to complete multiple projects, subject to a fixed budget.

Different from all the above works where they aim to select a best qualified team, our objective is to group individuals into multiple teams. It turns out that these two problems are closely related, this allows us to leverage existing techniques for team formation to solve our problem.

The other category of related work is maximum disjoint set cover problem (Bagaria et al. 2013). Given a universe, and a set of subsets, the objective is to find as many set covers as possible such that all set covers are pairwise disjoint. Our problem can be considered as a generalized disjoint set cover problem in the sense that every task in our problem may have different requirement of coverage, capacity constraint, and profit, and any feasible set cover must satisfy both coverage requirement as well as social compatibility. In addition, the requirement of "disjoint" is also relaxed by allowing individuals to have different load limits in our problem. Therefore, this work contributes fundamentally to the generalized cover decomposition problem.

3. Problem Formulation

Individuals. Skills. Tasks. In this paper we will assume that there is a set of n skills $\mathcal{S}$, a set of m individuals $\mathcal{V}$ and a set of k tasks $\mathcal{T}$. Each individual $u \in \mathcal{V}$, is represented by a subset of skills, i.e., $u \subseteq \mathcal{S}$; these are the skills that the individual possesses. Similarly, every task $t \in \mathcal{T}$ can also be represented by the set of skills that are required in order for the task to be completed (i.e., $t \subseteq \mathcal{S}$). In addition, each task t is associated with a profit λ_t upon the completion of this task. We assume that each task has unlimited number of copies, i.e., the same task can be performed by multiple teams. Notice that this assumption may not always hold in real world, to this end, we also study the case where each task has a capacity constraint, i.e., task t can only be performed up to g_t times.

Load. Our basic model assumes that each individual can only participate in *one* task. In Section 5, we will relax this assumption by allowing individuals to have different load limits, i.e.,, each individual u can participate in up to f_u number of tasks.

Teams. In practice, the social compatibility among individuals play an important role in a team work. For example, low social compatibility or high coordination cost will degrade the efficiency of organizations (Coase 1937). We model the *social compatibility* by means of a social network $G = (\mathcal{V}, \mathcal{E})$. One natural and popular option with respect to capturing the underlying social compatibility of a team is *connectivity*. This follows the approach

of (Lappas et al. 2009) and requires that each team C forms a connected graph. It was worth noting that there exist many ways to quantify the social compatibility among individuals, other options include the *diameter constraint* (Anagnostopoulos et al. 2012), i.e., the longest shortest path among team members in G is no larger than a given threshold. Fortunately, our results are not restricted to any specific measures of social compatibility, instead, we propose a general framework that works for any measure of social compatibility that has been explicitly defined.

Problem Formulation. For a team of individuals $C \subseteq \mathcal{V}$, we say that team C has a skill s if there exist at least one individual $u \in C$, such that u has skill s, i.e., $s \in u$. For a task $t \in \mathcal{T}$, we say that team C covers t if C (as a team) has all the skills required for t. Clearly, a team of individuals may cover more than one tasks, but they can only participate in one of those tasks due to each individual's load limit[1]. We define the set of qualified teams for task t to be the set of distinct teams that is social compatible and covers task t. That is,

$$\mathcal{C}_t = \{C \subseteq \mathcal{V} | C \text{ is social compatable and covers } t\}$$

A minimal qualified team of a task t is a qualified team of this task that is not a superset of any other qualified team. In the rest of this paper, we only consider minimal qualified teams and let $\mathbf{C} = \{\mathcal{C}_1, \cdots, \mathcal{C}_k\}$ denote the set of sets of minimal qualified teams for all tasks. The objective of this work is to find a most profitable way to group individuals into different teams, and assign one task to each team, such that (i) each team possesses all skills required by the corresponding task, (ii) all team members are social compatible, and (iii) each individual can only participate in one team. Given the above notation and constraint, we can now define TeamGrouping problem as follows:

P.1: *Maximize* $\sum_{C_{ti} \in \mathcal{C}_t \in \mathbf{C}} (x_{ti} \cdot \lambda_t)$

subject to:

$$\begin{cases} \sum_{u \in C_{ti} \in \mathcal{C}_t \in \mathbf{C}} x_{ti} \leq 1, \forall u \in \mathcal{V} \\ x_{ti} \in \{0,1\}, \forall C_{ti} \in \mathcal{C}_t \in \mathbf{C} \end{cases}$$

In the above formulation, x_{ti} indicates whether team C_{ti} has been selected ($x_{ti} = 1$) or not ($x_{ti} = 0$), and the first constraint specifies the load limit on each individual. The following results show that we cannot hope to achieve an $\omega(1/\ln m)$ approximation ratio for this problem.

[1] As mentioned earlier,this assumption will be relaxed in Section 5.

THEOREM 1. *The* TEAMGROUPING *problem is* $1/\ln m$ *hard to approximate.*

Proof: For our proof, we will consider a simplified version of TEAMGROUPING problem with only one task, i.e., $k=1$, and there is no constraint on social compatibility. We call this problem S-TEAMGROUPING. We next prove that the *maximum disjoint set cover cover problem* (DSCP) can be reduced to S-TEAMGROUPING. The formal definition of DSCP is as follows: Given a universe $\mathcal{U}$, and a set of subsets $\mathcal{X}$, find as many set covers as possible such that all set covers are pairwise disjoint. We wish to formulate an equivalent S-TEAMGROUPING with a set of skills $\mathcal{S}$ required the task, and a set of individuals $\mathcal{V}$. Let $\mathcal{S}=\mathcal{U}$ and $\mathcal{V}=\mathcal{X}$. Because there is only one task and no constraint on social compatibility, S-TEAMGROUPING is equivalent to grouping $\mathcal{V}$ into maximum number of disjoint teams each of which can cover all skills in $\mathcal{S}$. It was shown in (Bagaria et al. 2013) that the DSCP is hard to achieve an $\omega(1/\ln m)$ approximation ratio unless $NP \subseteq DTIME(n^{O(\ln\ln m)})$, thus TEAMGROUPING, which is a general case of S-TEAMGROUPING, is also $1/\ln m$ hard to approximate. □

One immediate result from the above proof is that if there is only one task and no constraint on social compatibility, we can simply adopt the method proposed in (Bagaria et al. 2013) to achieve $1/\ln m$ approximation ratio. In the following, we propose a LP-based approximation algorithm to tackle the general case.

4. LP-Based Approximation Algorithms

In this section, we give a $\max\{\mu/\Delta, \mu/2\sqrt{m}\}$-approximation algorithm for TEAMGROUPING, where μ is the approximation factor of the algorithm for the MINCOSTTEAMSELECTION problem, and $\Delta = \max_{C\in\mathcal{C}_t\in\mathbf{C}}|C|$, i.e., the size of the largest minimal team. The formal definition of MINCOSTTEAMSELECTION will be introduced in Definition 1.

4.1. LP Relaxation

Primal LP of P.1: *Maximize* $\sum_{C_{ti}\in\mathcal{C}_t\in\mathbf{C}}(x_{ti}\cdot\lambda_t)$

subject to:

$$\begin{cases}\sum_{u\in C_{ti}\in\mathcal{C}_t\in\mathbf{C}} x_{ti}\le 1, \forall u\in\mathcal{V}\\ 0\le x_{ti}\le 1, \forall C_{ti}\in\mathcal{C}_t\in\mathbf{C}\end{cases}$$

The above is the linear program (LP) relaxation of **P.1**. This LP has m constraints (excluding the trivial constraints $x_{ti}\ge 0, \forall C_{ti}\in\mathcal{C}_t\in\mathbf{C}$). However, since the number of variables $\sum_{t\in\mathcal{T}}|\mathcal{C}_t|$ could easily be exponential in the number of individuals, standard LP solvers can not solve this packing LP effectively.

To tackle this challenge, we adopt ellipsoid algorithm (Grötschel et al. 1981) which is capable of solving certain LP problems where the number of constraints is exponential in polynomial time. The idea of the ellipsoid method can be roughly described in the following. Given a non-degenerate convex set S, we would like to test whether S is empty or not. We start with an ellipsoid which is guaranteed to contain S. At each iteration, we check if the center of the current ellipsoid is in S. If it is, we are done, we can conclude that S is nonempty. Otherwise, we take a hyperplane through the center such that S is contained in of the two half-ellipsoids. We take the smallest ellipsoid completely containing this half-ellipsoid, whose volume is substantially smaller than the volume of the previous ellipsoid. We iterate on this new ellipsoid. In the worst case we need to iterate until the volume of the bounding ellipsoid gets below a pre-specified threshold value, in which case we can conclude that S is empty. It turns out that only a polynomial number of iterations are required in the case of linear programming. The algorithm does not require an explicit description of the linear program. All that is required is a polynomial time Separating Oracle, which checks whether a point lies in S or not, and returns a separating hyperplane in the latter case.

We refer to the above relaxed TeamGrouping problem as the primal LP. The dual to this primal LP associates a price $y(u)$ for each node $u \in \mathcal{V}$:

Dual LP of P.1: *Minimize* $\sum_{u\in\mathcal{V}} y(u)$

subject to:

$$\begin{cases} \sum_{u\in C_{ti}} y(u) \geq \lambda_t, \forall C_{ti} \in \mathcal{C}_t \in \mathbf{C} \\ y(u) \geq 0, \forall u \in \mathcal{V} \end{cases}$$

We leverage the ellipsoid method for exponential-sized LP with an (approximate) separation oracle to establish an approximation-preserving reduction from MincostTeamSelection, as defined in the following, to primal LP.

Definition 1 (MincostTeamSelection). Assume that there is a set of skills $\mathcal{S}$ and individuals $\mathcal{V}$, each individual $u \in \mathcal{V}$ is associated with a cost and possesses a subset of skills. Find a team of individuals with minimum cost such that (1) all team members are social compatible, and (2) all skills in $\mathcal{S}$ can be covered.

Depending on the definition of social compatibility, MincostTeamSelection has been intensively studied in the literature. In (Lappas et al. 2009), they propose to use connectivity as a measure of social compatibility, that is, all team members must be connected

in the social network. Under this context, the MincostTeamSelection problem can be reduced from *node weight group steiner tree* problem (Khandekar et al. 2012) which admits a performance ratio of $O(|\mathcal{E}|^{1/2}\ln|\mathcal{E}|)$ where $|\mathcal{E}|$ is the number of edges in the social network. It was worth noting that condition (1) can be replaced by other reasonable measurements on social compatibility among team members, for instance, some work (Anagnostopoulos et al. 2012) requires that a team must have bounded diameter. The following theorem is not restricted to any specific measure of social compatibility.

Theorem 2. *If there is a polynomial μ-approximation algorithm for* MincostTeamSelection*, then there exists a polynomial μ-approximation algorithm for* ***P.1****.*

Proof: Let $\mathcal{A}$ be a μ-approximation algorithm for MincostTeamSelection. We run the ellipsoid algorithm on the dual LP using the algorithm $\mathcal{A}$ as the approximate separation oracle. More precisely, let $S(L)$ denote the set of $y \in \mathbb{R}_+^{\mathcal{V}}$ satisfying that

$$\sum_{u\in\mathcal{V}} y(u) \le L$$

$$\sum_{u\in C_{ti}} y(u) \ge \lambda_t, \forall C_{ti} \in \mathcal{C}_t \in \mathbf{C}$$

We can adopt binary search to find the smallest value of L for which $S(L)$ is non-empty. The separation oracle works as follows: First, it checks the inequality $\sum_{u\in\mathcal{V}} y(u) \le L$. Next, it runs the algorithm $\mathcal{A}$ on each task $t \in \mathcal{T}$ and select a group $C_{ti} \in \mathcal{C}_t, \forall t \in \mathcal{T}$, using $y(u)$ as the price function. If for all t and i, C_{ti} has cost larger than λ_t, then we know that $y \in S(L)$. Otherwise, if there exists some t and i such that C_{ti} has cost less than λ_t, then we conclude that $y \in S(L)$ and C_{ti} gives us a separating hyperplane. However, since $\mathcal{A}$ is just an approximation algorithm, the above conclusion might be incorrect, i.e., $S(L)$ might actually be empty. Fortunately, since the approximation factor of $\mathcal{A}$ is at most μ, we know that in this case, $\mu \cdot y \in S(u \cdot L)$. Therefore, let L^* be the minimum value of L for which the algorithm decides $S(L)$ is non-empty, then we know that $S(L^* - \epsilon)$ is empty (where ϵ depends on the precision of the algorithm), and $S(\mu L^*)$ is nonempty. Therefore, the value of the dual LP, and hence, the value of the primal LP, is between L^* and μL^*.

The above algorithm computes the approximate value of the primal LP. Next, we describe how to compute the actual approximate solution. Now, let $\mathcal{C}_t^H$ denote the subset of teams on task t corresponding to the separating hyper-planes found by the above separation

oracle while running the ellipsoid algorithm on $S(L^* - \epsilon)$. Then, $\sum_{t=1}^{k} |\mathcal{C}_t^H|$ is polynomial. Define $\mathbf{C}^H = \{\mathcal{C}_1^H, \cdots, \mathcal{C}_k^H\}$, consider the restricted dual LP:

$$\begin{aligned}
&\textit{Minimize } \textstyle\sum_{u \in \mathcal{V}} y(u) \\
&\textbf{subject to:} \\
&\qquad \begin{cases} \sum_{u \in C_{ti}} y(u) \geq \lambda_t, \forall C_{ti} \in \mathcal{C}_t^H \in \mathbf{C}^H \\ y(u) \geq 0, \forall u \in \mathcal{V} \end{cases}
\end{aligned}$$

Its value is also at least L^*. So, we solve the following restricted primal LP of polynomial size, which is the dual of the restricted dual LP:

$$\begin{aligned}
&\textit{Maximize } \textstyle\sum_{C_{ti} \in \mathcal{C}_t^H \in \mathbf{C}^H} (x_{ti} \cdot \lambda_t) \\
&\textbf{subject to:} \\
&\qquad \begin{cases} \sum_{u \in C_{ti} \in \mathcal{C}_t^H \in \mathbf{C}^H} x_{ti} \leq 1, \forall u \in \mathcal{V} \\ 0 \leq x_{ti} \leq 1, \forall C_{ti} \in \mathcal{C}_t^H \in \mathbf{C}^H \end{cases}
\end{aligned}$$

The optimal solution of this restricted LP has value at least L^*, which is a μ-approximation to the original primal LP. □

4.2. Approximation Algorithm

Having described the LP relaxation, we now propose an approximation algorithms computing a group of teams from LP solutions. Our approach involves two algorithms as subroutines.

Candidate Solution I: In the first algorithm (Algorithm 1), we directly apply the deterministic rounding (Algorithm 5) to $\mathbf{C}^H$. We can prove that this algorithm achieves μ/Δ approximation ratio where Δ denotes the size of the largest minimal team. For ease of presentation, we put the detailed description of our rounding technique in Section 4.2.1.

Algorithm 1 Candidate Grouping - I

1: Apply deterministic rounding (Algorithm 5) to $\mathbf{C}^H$ and output a group of teams.

LEMMA 1. *Algorithm 1 achieves μ/Δ approximation ratio for* TEAMGROUPING.

Proof: In Theorem 6, we prove that our deterministic rounding technique achieves $1/\rho$ approximation ratio where ρ is the size of the largest possible team in $\mathbf{C}^H$. We know that $\rho \leq \Delta$, therefore the profit of the solution returned from Algorithm 1 is at least

$$\frac{1}{\Delta} \sum_{C_{ti} \in \mathcal{C}_t^H \in \mathbf{C}^H} (x_{ti}^* \cdot \lambda_t) \geq \frac{\mu}{\Delta} \cdot OPT$$

where x^*_{ti} is the solution of primal LP and OPT denotes the amount of profits gained from the optimal grouping. □

Candidate Solution II: The framework of the second candidate solution (Algorithm 2) can be summarized as follows:

Step 1: Recall that $\mathcal{C}^H_t$ denotes the subset of teams on task t corresponding to the separating hyper-planes found by primal LP. For every task t, we first partition $\mathcal{C}^H_t$ to two disjoint subsets $\mathcal{C}^{H_1}_t$ and $\mathcal{C}^{H_2}_t$ such that: $\forall C \in \mathcal{C}^{H_1}_t : |C| \leq \rho$ and $\forall C \in \mathcal{C}^{H_2}_t : |C| > \rho$. That is, $\mathcal{C}^{H_1}_t$ contains all teams with less than ρ individuals. Let $\mathbf{C}^{H_1} = \{\mathcal{C}^{H_1}_1, \cdots, \mathcal{C}^{H_1}_k\}$ and $\mathbf{C}^{H_2} = \{\mathcal{C}^{H_2}_1, \cdots, \mathcal{C}^{H_2}_k\}$.

Step 2: Apply deterministic rounding (Algorithm 5) to $\mathbf{C}^{H_1}$ and output a group of teams $\widetilde{\mathcal{C}}$.

Step 3: Select a team, say $C_{t_{\max}}$, from $\mathbf{C}^{H_2}$ whose task $t_{\max}$ has the highest profit $\lambda_{t_{\max}}$.

Step 4: Compare $\widetilde{\mathcal{C}}$ and $\{C_{t_{\max}}\}$, choose the one with larger profit as the final output, i.e, the profit of the returned solution is $\max\{\sum_{C_{ti}\in\widetilde{\mathcal{C}}} \lambda_t, \lambda_{t_{\max}}\}$.

Algorithm 2 Candidate Grouping - II

1: Partition $\mathbf{C}^H$ into two subsets $\mathbf{C}^{H_1}$ and $\mathbf{C}^{H_2}$
2: Apply the deterministic rounding (Algorithm 5) to $\mathbf{C}^{H_1}$ and output $\widetilde{\mathcal{C}}$.
3: Select a team with the highest profit, say $C_{t_{\max}}$, from $\mathbf{C}^{H_2}$.
4: Compare $\widetilde{\mathcal{C}}$ and $\{C_{t_{\max}}\}$, return the one with larger profit.

We next prove that the approximation ratio of Algorithm 2 can be bounded by $\mu/2\sqrt{m}$.

LEMMA 2. *Algorithm 2 achieves $\mu/2\sqrt{m}$ approximation ratio for* TEAMGROUPING.

Proof: To prove this lemma, it suffices to show that $\max\{\sum_{C_{ti}\in\widetilde{\mathcal{C}}} \lambda_t, \lambda_{t_{\max}}\} \geq \frac{1}{\sqrt{m}} \cdot \frac{\mu}{2} \cdot OPT$.

We first bound the gap between the profit gained from $\widetilde{\mathcal{C}}$ and $\sum_{C_{ti}\in\mathcal{C}^{H_1}_t\in\mathbf{C}^{H_1}} (x^*_{ti} \cdot \lambda_t)$. As proved in Lemma 6, our deterministic rounding (Algorithm 5) achieves approximation ratio $1/\rho$. Because $\rho \leq \sqrt{m}$ holds for all teams from $\mathbf{C}^{H_1}$, we have

$$\sum_{C_{ti}\in\widetilde{\mathcal{C}}} \lambda_t \geq \frac{1}{\rho} \cdot \sum_{C_{ti}\in\mathcal{C}^{H_1}_t\in\mathbf{C}^{H_1}} (x^*_{ti} \cdot \lambda_t) \geq \frac{1}{\sqrt{m}} \cdot \sum_{C_{ti}\in\mathcal{C}^{H_1}_t\in\mathbf{C}^{H_1}} (x^*_{ti} \cdot \lambda_t) \tag{1}$$

We next bound the gap between the profit gained from $C_{t_{\max}}$ and $\sum_{C_{ti}\in\mathcal{C}^{H_2}_t\in\mathbf{C}^{H_2}} (x^*_{ti} \cdot \lambda_t)$. In particular, we show that $\lambda_{t_{\max}} \geq \frac{1}{\sqrt{m}} \cdot \sum_{C_{ti}\in\mathcal{C}^{H_2}_t\in\mathbf{C}^{H_2}} (x^*_{ti} \cdot \lambda_t)$.

First, we have

$$\sum_{C_{ti}\in\mathcal{C}_t^{H_2}\in\mathbf{C}^{H_2}}(x_{ti}^*\cdot\lambda_t)\leq\sum_{C_{ti}\in\mathcal{C}_t^{H_2}\in\mathbf{C}^{H_2}}(x_{ti}^*\cdot\lambda_{t_{\max}})=\lambda_{t_{\max}}\cdot\sum_{C_{ti}\in\mathcal{C}_t^{H_2}\in\mathbf{C}^{H_2}}x_{ti}^*\leq\lambda_{t_{\max}}\cdot\frac{m}{\sqrt{m}}\qquad(2)$$

The first inequality comes from the assumption that $C_{t_{\max}}$ delivers the highest profit among $\mathbf{C}^{H_2}$. The last inequality is due to the following observation: because the load of each individual is at most 1, we have $\sum_{C_{ti}\in\mathcal{C}_t^{H_2}\in\mathbf{C}^{H_2}}(x_{ti}^*\cdot|C_{ti}|)\leq m$, recall that all teams in $\mathbf{C}^{H_2}$ contain at least $\sqrt{m}$ individuals, we have $\sum_{C_{ti}\in\mathcal{C}_t^{H_2}\in\mathbf{C}^{H_2}}(x_{ti}^*\cdot|C_{ti}|)\geq\sum_{C_{ti}\in\mathcal{C}_t^{H_2}\in\mathbf{C}^{H_2}}(x_{ti}^*\cdot\sqrt{m})$, then we have $\sum_{C_{ti}\in\mathcal{C}_t^{H_2}\in\mathbf{C}^{H_2}}(x_{ti}^*\cdot\sqrt{m})\leq m$, thus $\sum_{C_{ti}\in\mathcal{C}_t^{H_2}\in\mathbf{C}^{H_2}}(x_{ti}^*)\leq\sqrt{m}$.

In addition, based on Theorem 2, we have

$$\sum_{C_{ti}\in\mathcal{C}_t^{H}\in\mathbf{C}^{H}}(x_{ti}^*\cdot\lambda_t)\geq\mu\cdot OPT\Rightarrow\sum_{C_{ti}\in\mathcal{C}_t^{H_1}\in\mathbf{C}^{H_1}}(x_{ti}^*\cdot\lambda_t)+\sum_{C_{ti}\in\mathcal{C}_t^{H_2}\in\mathbf{C}^{H_2}}(x_{ti}^*\cdot\lambda_t)\geq\mu\cdot OPT\qquad(3)$$

Eqs. (1) (2) (3) together imply that

$$\max\{\sum_{C_{ti}\in\widetilde{\mathcal{C}}}\lambda_t,\lambda_{t_{\max}}\}\geq\frac{1}{\sqrt{m}}\cdot\frac{\sum_{C_{ti}\in\mathcal{C}_t^{H_1}\in\mathbf{C}^{H_1}}(x_{ti}^*\cdot\lambda_t)+\sum_{C_{ti}\in\mathcal{C}_t^{H_2}\in\mathbf{C}^{H_2}}(x_{ti}^*\cdot\lambda_t)}{2}\geq\frac{1}{\sqrt{m}}\cdot\frac{\mu}{2}\cdot OPT$$

□

Algorithm 3 Approx-TG

1: Compute two candidate solutions using Algorithm 1 and Algorithm 2.
2: Return the one with higher profit.

Putting It All Together. Given solutions returned from Algorithm 1 and Algorithm 2, we simply choose the one with higher profit as our final output. We refer to this algorithm as Approx-TG (Algorithm 3). Lemma 1 and Lemma 2 together imply our main theorem.

THEOREM 3. *Approx-TG achieves* $\max\{\mu/\Delta,\mu/2\sqrt{m}\}$ *approximation ratio for* TEAMGROUPING.

Now consider a special case of TEAMGROUPING where there is no constraint on social compatibility. Under this setting, MINCOSTTEAMSELECTION problem as defined in Definition 1 is equivalent to classic *weighted set cover problem* (Chvatal 1979), which allows $\ln n$ approximation. In addition, we have $\Delta\leq n$, this is because the number of possible skills is at most n, if there is no constraint on social compatibility, any minimal qualified team contains at most n individuals. Then the following corollary holds by replacing μ using $\ln n$, and Δ using n in Theorem 3.

COROLLARY 1. *If there is no constraint on social compatibility, Approx-TG achieves* $\max\{\ln n/n, \ln n/2\sqrt{m}\}$ *approximation ratio for* TEAMGROUPING.

It was worth noting that in practise, $n \ll m$, i.e, the number of skills is much smaller than the number of individuals, thus the above approximation ratio can be further rewritten as $\ln n/n$.

Consider another special case that uses connectivity to measure the social compatibility. As discussed earlier in Section 4.1, under this setting, the MINCOSTTEAMSELECTION problem can be reduced from *node weight group steiner tree* problem (Khandekar et al. 2012) which admits a performance ratio of $O(|\mathcal{E}|^{1/2} \ln |\mathcal{E}|)$. Therefore, we have the following corollary.

COROLLARY 2. *If all teams are required to be connected, Approx-TG achieves* $\max\{O(|\mathcal{E}|^{1/2} \ln |\mathcal{E}|)/\Delta), O(|\mathcal{E}|^{1/2} \ln |\mathcal{E}|)/2\sqrt{m})\}$ *approximation ratio for* TEAMGROUPING.

4.2.1. LP Rounding We next discuss how to round the factional solution of primal LP. We propose two rounding techniques: *randomized rounding* and *deterministic rounding*. Both techniques can be used in Algorithm 3 as subroutines. In the rest of our discussion, we say two teams are *adjacent* if they contain at least one common individual. We use $\mathcal{N}(C)$ to denote the adjacent teams of C. Let $\mathbf{C}^I$ denote the set of input teams, e.g., $\mathbf{C}^I$ refers to $\mathbf{C}^H$ (or $\mathbf{C}^{H_1}$ resp.) in Algorithm 1 (or Algorithm 2 resp.).

Randomized Rounding Our randomized rounding (Algorithm 4) consists of two major parts: a rounding stage and a conflict resolution stage. In the first stage, a initial team set is generated as follows. For each team C_{ti} under consideration, the status (whether is removed or not) is determined independently at random. Recall that x^*_{ti} denotes the solution of primal LP. Let ρ denote the size of the largest team in $\mathbf{C}^I$, each team is taken with probability $\frac{x^*_{ti}}{2\rho}$ to survive and with the remaining probability to be removed. Conflicts can occur when two adjacent teams both survive. In this case, assume each team has a unique index, the conflict is resolved by letting the team with smaller index survive. The other team is removed from the solution by being allocated the empty set.

LEMMA 3. *Resulted from any feasible LP solution* x^*_{ti}*, team* C_{ti} *survives after the first phase with probability* $\frac{x^*_{ti}}{2\rho}$.

Algorithm 4 Randomized Rounding

1: **for** each team C_{ti} in $\mathbf{C}^I$ **do**
2: Let C_{ti} survive with probability $\frac{x^*_{ti}}{2\rho}$
3: **for** each C_{ti} that is survived from last phase **do**
4: Remove C_{ti} if some team in $\mathcal{N}(C_{ti})$ with smaller index also survives
5: Return the remaining teams as output.

The above lemma can be directly derived from our algorithm description. Next we use 0/1 random variable X_{ti} to denote whether x_{ti} is set to 1 after the first phase, we can immediately have $\mathbf{E}[X_{ti}] = \frac{x^*_{ti}}{2\rho}$. Let Y_{ti} be a 0/1 random variable representing whether x_{ti} is set to 1 after the second phase. The event that C_{ti} survives at the first phase but removed from the second phase can be represented as: $Y_{ti} = 0$, under the condition that $X_{ti} = 1$. And the probability of this event is $\Pr[Y_{ti} = 0 | X_{ti} = 1]$.

LEMMA 4. *For any team C_{ti} that is having survived in the first phase, the probability that C_{ti} still survives after the second phase is at least $\frac{1}{2}$.*

Proof: We note that this event can only happen if $X_{lj} = 1$ for some $C_{lj} \in \mathcal{N}(C_{ti})$, that is

$$\sum_{C_{lj} \in \mathcal{N}(C_{ti})} X_{lj} \geq 1$$

Then together with Markov's inequality, the probability of this event can be bounded by

$$\Pr[Y_{ti} = 0 | X_{ti} = 1] \leq \Pr[\sum_{C_{lj} \in \mathcal{N}(C_{ti})} X_{lj} \geq 1] \leq \mathbf{E}[\sum_{C_{lj} \in \mathcal{N}(C_{ti})} X_{lj}]$$

Based on linearity of expectation and $\mathbf{E}[X_{lj}] = \frac{x^*_{ti}}{2\rho}$, we have:

$$\mathbf{E}[\sum_{C_{lj} \in \mathcal{N}(C_{ti})} X_j] = \sum_{C_{lj} \in \mathcal{N}(C_{ti})} \mathbf{E}[X_{lj}] = \sum_{C_{lj} \in \mathcal{N}(C_{ti})} \frac{x^*_{lj}}{2\rho}$$

Recall that all x^*_{lj} satisfy LP constraints and each team has at most ρ individuals, then according to first constraint in primal LP, we further have:

$$\sum_{C_{lj} \in \mathcal{N}(C_{ti})} \frac{x^*_{lj}}{2\rho} = \frac{\sum_{C_{lj} \in \mathcal{N}(C_{ti})} x^*_{lj}}{2\rho} \leq \frac{1 \cdot \rho}{2\rho} = \frac{1}{2}$$

Therefore each team still survives after the second phase is at least $1 - \frac{1}{2} = \frac{1}{2}$. □

Lemma 3 and Lemma 2 together imply that each team C_{ti} survives with probability $x^*_{ti}/4\rho$, then the following theorem follows.

LEMMA 5. *Algorithm 4 achieves approximation ratio $1/4\rho$.*

Deterministic Rounding We next propose a deterministic rounding method (Algorithm 5) with approximation ratio $1/\rho$. Our method can be described as follows:

Step 1: Sort all teams in $\mathbf{C}^I$ in non-decreasing order of their profit.

Step 2: Select the team $C_{ti} \in \mathbf{C}^I$ with the highest profit and add it to our final solution.

Step 3: Remove C_{ti} and $\mathcal{N}(C_{ti})$ from $\mathbf{C}^I$. This step ensures that no individual participates in multiple tasks.

Step 4: Goto Step 2 unless there are no teams left.

Algorithm 5 Deterministic Rounding

1: Sort all teams in $\mathbf{C}^I$ in non-decreasing order of their profit.
2: **while** $\mathbf{C}^I \neq \emptyset$ **do**
3: Select the team with highest profit in $\mathbf{C}^I$, say C_{ti}
4: $\mathcal{C}^{DR} = \mathcal{C}^{DR} \cup \{C_{ti}\}$
5: $\mathbf{C}^I = \mathbf{C}^I \setminus \{C_{ti} \cup \mathcal{N}(C_{ti})\}$
6: Return $\mathcal{C}^{DR}$

We next provide the performance analysis of Algorithm 5.

LEMMA 6. *Algorithm 5 achieves approximation ratio* $1/\rho$.

Proof: Let $\mathcal{C}^{DR}$ denote the group of teams selected by Algorithm 5. Consider any team $C_{ti} \in \mathcal{C}^{DR}$, we have $x^*_{ti} \cdot \lambda_t + \sum_{C_{lj} \in \mathcal{N}(C_{ti})} (x^*_{lj} \cdot \lambda_l) \leq \rho \cdot \lambda_t$. This is because C_{ti} has the highest profit among all its adjacent teams and $\sum_{u \in C_{lj} \in \mathcal{C}_l \in \mathbf{C}} x_{ti} \leq 1, \forall u \in C_{ti}$. Then the following holds for any $C_{ti} \in \mathcal{C}^{DR}$: $\lambda_t \geq (x^*_{ti} \cdot \lambda_t + \sum_{C_{lj} \in \mathcal{N}(C_{ti})} (x^*_{lj} \cdot \lambda_l))/\rho$, this implies that the total profit of $\mathcal{C}^{DR}$ is at least $1/\rho$ of the optimal solution. □

5. Extensions

5.1. Incorporating the Heterogeneity of Each Individual's Load Limit

Our basic model assumes that each individual can only participate in *one* task. However, as mentioned earlier, different individuals may have different capabilities, i.e., , each individual u can participate in up to f_u number of tasks. In order to capture this scenario, we can simply create f_u copies of u with identical skill set, then all results developed previously can apply to the modified instance.

THEOREM 4. *Algorithm achieves* $\max\{\mu/\Delta, \mu/2\sqrt{m}\}$ *approximation ratio when individuals have heterogeneous load limits.*

5.2. Incorporating the Capacity Constraint of Each Task

Throughout this paper, we assume that each task can be performed unlimited number of times. However, this may not always hold in practice, take puzzle assembly as an example, this type of task can only be performed once. To this end, we add a group of additional constraints to the original problem: $\sum_{C_{ti}\in\mathcal{C}_t} x_{ti} \leq g_t, \forall t \in \mathcal{T}$ where g_t denotes the capacity of task $t \in \mathcal{T}$, i.e., task t can be performed up to g_t times.

P.2: *Maximize* $\sum_{C_{ti}\in\mathcal{C}_t\in\mathbf{C}}(x_{ti}\cdot\lambda_t)$

subject to:

$$\begin{cases}\sum_{u\in C_{ti}\in\mathcal{C}_t\in\mathbf{C}} x_{ti} \leq 1, \forall u \in \mathcal{V} \\ \sum_{C_{ti}\in\mathcal{C}_t} x_{ti} \leq g_t, \forall t \in \mathcal{T} \\ x_{ti} \in \{0,1\}, \forall C_{ti} \in \mathcal{C}_t \in \mathbf{C}\end{cases}$$

Similar to the one developed in Section 4, we propose a LP-Based Approximation Algorithm for **P.2**.

LP Relaxation The primal LP of **P.2** can be formulated as follows.

Primal LP of P.2: *Maximize* $\sum_{C_{ti}\in\mathcal{C}_t\in\mathbf{C}}(x_{ti}\cdot\lambda_t)$

subject to:

$$\begin{cases}\sum_{u\in C_{ti}\in\mathcal{C}_t\in\mathbf{C}} x_{ti} \leq 1, \forall u \in \mathcal{V} \\ \sum_{C_{ti}\in\mathcal{C}_t} x_{ti} \leq g_t, \forall t \in \mathcal{T} \\ 0 \leq x_{ti} \leq 1, \forall C_{ti} \in \mathcal{C}_t \in \mathbf{C}\end{cases}$$

The dual to the above primal LP associates a price $y(u)$ for each node $u \in V$ and a price $p(t)$ for each task $t \in \mathcal{T}$:

Dual LP of P.2: *Minimize* $\sum_{u\in\mathcal{V}} y(u) + \sum_{t\in\mathcal{T}}(p(t)\cdot g_t)$

subject to:

$$\begin{cases}\sum_{u\in C_{ti}} y(u) + p(t) \geq \lambda_t, \forall C_{ti} \in \mathcal{C}_t \in \mathbf{C} \\ y(u) \geq 0, \forall u \in \mathcal{V}; p(t) \geq 0, \forall t \in \mathcal{T}\end{cases}$$

Similar to the solution for **P.1**, we run the ellipsoid algorithm on the dual LP using the algorithm $\mathcal{A}$ as the approximate separation oracle. More precisely, let $S(L)$ denote the set of $y \in \mathbb{R}_+^{\mathcal{V}}$ satisfying that

$$\sum_{u\in\mathcal{V}} y(u) + \sum_{t\in\mathcal{T}}(p(t)\cdot g_t) \leq L$$

$$\sum_{u\in C_{ti}} y(u) + p(t) \geq \lambda_t, \forall C_{ti} \in \mathcal{C}_t \in \mathbf{C}$$

We adopt binary search to find the smallest value of L for which $S(L)$ is non-empty. The separation oracle works as follows: First, it checks the inequality $\sum_{u\in\mathcal{V}} y(u) + \sum_{t\in\mathcal{T}}(p(t)\cdot$

$g_t) \leq L$. Next, it runs the algorithm $\mathcal{A}$ on each task $t \in \mathcal{T}$ and select a group $C_{ti} \in \mathcal{C}_t, \forall t \in \mathcal{T}$, using $y(u)$ as the price function. If for all t and i, C_{ti} has cost larger than $\lambda_t - p(t)$, then we know that $y \in S(L)$. Otherwise, if there exists some t and i such that C_{ti} has cost less than $\lambda_t - p(t)$, then we conclude that $y \in S(L)$ and C_{ti} gives us a separating hyperplane. Then based on similar analysis in Section 4, we have the following theorem.

THEOREM 5. *If there is a polynomial μ-approximation algorithm for* MINCOSTTEAMSELECTION*, then there exists a polynomial μ-approximation algorithm for* ***P.2****.*

LP Rounding This stage is different from the one used in Section 4.2.1. In Section 4.2.1, we only consider the load limit on each individual, therefore, it suffice to ensure no adjacent teams are selected at the same time. When taking into account the capacity constraint, one additional requirement is not to assign too many teams to the same task. Take the deterministic rounding technique as an example, our method (Algorithm 6) can be described as follows:

Step 1: Sort all teams in $\mathbf{C}^I$ in non-decreasing order of their profit.

Step 2: Select the team C_{ti} with the highest profit and set $x^*_{ti} = 1$.

Step 2.1: Let $\mathcal{C}^I_t \in \mathbf{C}^I$ denote all teams assigned to task t (with positive x^*_{tj}) in $\mathbf{C}^I$. Reduce the value of x^*_{tj} for some $C_{tj} \in \mathcal{C}^I_t \setminus \{C_{ti}\}$ to ensure that $\sum_{C_{tj} \in \mathcal{C}^I_t} x^*_{tj} \leq g_t$ after setting $x^*_{ti} = 1$. This can be done in many ways, one naive way is as follows: pick an arbitrary team in $C_{tj} \in \mathcal{C}^I_t \setminus \{C_{ti}\}$, reduce x^*_{tj} to some non-negative number (zero, if necessary), pick the next team and repeat the process, this process iterates team by team in an arbitrary order until $\sum_{C_{tj} \in \mathcal{C}^I_t} x^*_{tj} \leq g_t$ with $x^*_{ti} = 1$. This operation simply ensures that the capacity constraint on each task is satisfied after setting $x^*_{ti} = 1$.

Step 2.2: Remove $\mathcal{N}(C_{ti})$ and all teams in $\{C_{tj} | C_{tj} \in \mathcal{C}^I_t \setminus \{C_{ti}\} \bigwedge x^*_{tj} = 0\}$ from $\mathbf{C}^I$. This step ensures that no individual participates in multiple tasks and each task is assigned to at most g_t teams.

Step 3: Goto Step 2 unless there are no teams left.

We next analyze the approximation ratio of Algorithm 6.

LEMMA 7. *Algorithm 6 achieves approximation ratio $1/(\rho + 1)$.*

Proof: Let $\mathcal{C}^{DR}$ denote the group of teams selected by Algorithm 6. Consider any team $C_{ti} \in \mathcal{C}^{DR}$, because C_{ti} has the highest profit, we have

$$x^*_{ti} \cdot \lambda_t + \sum_{C_{lj} \in \mathcal{N}(C_{ti})} (x^*_{ti} \cdot \lambda_l) \leq \rho \cdot \lambda_t \tag{4}$$

Algorithm 6 Deterministic Rounding

1: Sort all teams in $\mathbf{C}^I$ in non-decreasing order of their profit.
2: **while** $\mathbf{C}^I \neq \emptyset$ **do**
3: Select the team with highest profit in $\mathbf{C}^I$, say C_{ti}
4: $\mathcal{C}^{DR} = \mathcal{C}^{DR} \cup \{C_{ti}\}$ and set $x^*_{ti} = 1$
5: Reduce the value of x^*_{tj} for some $C_{tj} \in \mathcal{C}^I_t \setminus \{C_{ti}\}$ to ensure that $\sum_{C_{tj} \in \mathcal{C}^I_t} x^*_{tj} \leq g_t$
6: Remove $\mathcal{N}(C_{ti})$ and all teams in $\{C_{tj} | C_{tj} \in \mathcal{C}^I_t \setminus \{C_{ti}\} \bigwedge x^*_{tj} = 0\}$ from $\mathbf{C}^I$
7: Return $\mathcal{C}^{DR}$

On the other hand, because all teams in $\mathcal{C}^I_t$ have equal profit λ_t, we have

$$x^*_{ti} \cdot \lambda_t + \delta \cdot \lambda_t \leq \lambda_t \tag{5}$$

where $\delta \leq 1 - x^*_{ti}$ denotes the reduced value of $\sum_{C_{tj} \in \mathcal{C}^I_t \setminus \{C_{ti}\}} x^*_{tj}$ in Step 2.2.

Eqs. (4) and (5) together imply that

$$(1+\rho)\lambda_t \geq \left(x^*_{ti} \cdot \lambda_t + \sum_{C_{lj} \in \mathcal{N}(C_{ti})} (x^*_{lj} \cdot \lambda_l) \right) + (x^*_{ti} \cdot \lambda_t + \delta \cdot \lambda_t) \geq x^*_{ti} \cdot \lambda_t + \sum_{C_{lj} \in \mathcal{N}(C_{ti})} (x^*_{lj} \cdot \lambda_l) + \delta \cdot \lambda_t$$

It follows that $\lambda_t \geq \frac{1}{1+\rho} \cdot (x^*_{ti} \cdot \lambda_t + \sum_{C_{lj} \in \mathcal{N}(C_{ti})} (x^*_{lj} \cdot \lambda_l) + \delta \cdot \lambda_t)$, indicating that we lose at most a factor of $\frac{\rho}{1+\rho}$ for the profit due to removing $\mathcal{N}(C_{ti})$ and all teams in $\{C_{tj} | C_{tj} \in \mathcal{C}^I_t \setminus \{C_{ti}\} \bigwedge x^*_{tj} = 0\}$ in Step 2.2. This finishes the proof. □

Approx-TG (Algorithm 3) can be naturally modified to handle this generalization by replacing its LP rounding by Algorithm 6. Following a similar proof of Theorem 3, we can prove the following theorem.

THEOREM 6. *The modified Approx-TG (Algorithm 3) achieves* $\max\{\mu/(\Delta + 1), \mu/2(\sqrt{m}+1)\}$ *approximation ratio for* ***P.2****.*

6. Conclusion

In this paper, we study the profit-aware team grouping problem. We assume a collection of tasks $\mathcal{T}$, where each task requires a specific set of skills, and yields a different profit upon completion. Active and qualified individuals may collaborate with each other in the form of *teams* to accomplish a set of tasks. We aim to group individuals into different teams, and assign them to different tasks, such that the total profit of the tasks that can be performed is maximized. We consider three constraints when perform grouping, and

present a LP-based approximation algorithm to tackle it. We also study several extensions of this problem. Although this paper studies team grouping problem, our results are general enough to tackle a broad range of generalized cover decomposition problems.

References

Anagnostopoulos, Aris, Luca Becchetti, Carlos Castillo, Aristides Gionis, Stefano Leonardi. 2012. Online team formation in social networks. *Proceedings of the 21st international conference on World Wide Web*. ACM, 839–848.

Bagaria, Vivek Kumar, Ashwin Pananjady, Rahul Vaze. 2013. Optimally approximating the coverage lifetime of wireless sensor networks. *arXiv preprint arXiv:1307.5230* .

Chvatal, Vasek. 1979. A greedy heuristic for the set-covering problem. *Mathematics of operations research* **4** 233–235.

Coase, Ronald H. 1937. The nature of the firm. *economica* **4** 386–405.

Dorn, Christoph, Schahram Dustdar. 2010. Composing near-optimal expert teams: A trade-off between skills and connectivity. *On the Move to Meaningful Internet Systems: OTM 2010*. Springer, 472–489.

Gajewar, Amita, Atish Das Sarma. 2012. Multi-skill collaborative teams based on densest subgraphs. *SDM*. SIAM, 165–176.

Golshan, Behzad, Theodoros Lappas, Evimaria Terzi. 2014. Profit-maximizing cluster hires. *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 1196–1205.

Grötschel, Martin, László Lovász, Alexander Schrijver. 1981. The ellipsoid method and its consequences in combinatorial optimization. *Combinatorica* **1** 169–197.

Kargar, Mehdi, Aijun An. 2011. Discovering top-k teams of experts with/without a leader in social networks. *Proceedings of the 20th ACM international conference on Information and knowledge management*. ACM, 985–994.

Kargar, Mehdi, Morteza Zihayat, Aijun An. 2013. Finding affordable and collaborative teams from a network of experts. *Proceedings of the SIAM International Conference on Data Mining (SDM)*. SIAM, 587–595.

Khandekar, Rohit, Guy Kortsarz, Zeev Nutov. 2012. Approximating fault-tolerant group-steiner problems. *Theoretical Computer Science* **416** 55–64.

Lappas, Theodoros, Kun Liu, Evimaria Terzi. 2009. Finding a team of experts in social networks. *Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 467–476.

Li, Cheng-Te, Man-Kwan Shan. 2010. Team formation for generalized tasks in expertise social networks. *Social Computing (SocialCom), 2010 IEEE Second International Conference on*. IEEE, 9–16.

Lu, Dawei. 2011. *Fundamentals of supply chain management*. Bookboon.

Pananjady, Ashwin, Vivek Kumar Bagaria, Rahul Vaze. 2014. Maximizing utility among selfish users in social groups. *Communications (NCC), 2014 Twentieth National Conference on*. IEEE, 1–6.

Sozio, Mauro, Aristides Gionis. 2010. The community-search problem and how to plan a successful cocktail party. *Proceedings of the 16th ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 939–948.